

A

JC945 U.S. DPTO

PTO
SEP 11 1961
AND

11/03/00

UTILITY PATENT APPLICATION TRANSMITTAL (Only for new nonprovisional applications under 37 CFR 1.53(b))	Attorney Docket No. First Inventor Title Express Mail Label No.	PD99-2788 Richard Willems SYSTEM AND METHOD FOR COLLECTING SYSTEM DATA USING RECORD BASED REQUESTS WITH TAG LISTS AND PAUSING ALL BUT ONE THREAD OF A COMPUTER SYSTEM EL700671685US
APPLICATION ELEMENTS		Assistant Commissioner for Patents Box Patent Application Washington, DC 20231
<div style="display: flex; justify-content: space-between;"> <div style="width: 48%;"> 1. <input checked="" type="checkbox"/> Fee Transmittal Form (submit an original and a duplicate for fee processing) 2. <input type="checkbox"/> Applicant claims small entity status. See 37 CFR 1.27 3. <input checked="" type="checkbox"/> Specification [total pages <u>39</u>] (preferred Arrangement set forth below) - Descriptive title of the Invention - Cross References to Related Applications - Statement Regarding Fed sponsored R&D - Reference to sequence listing, a table, or a computer program listing appendix - Background of the Invention - Brief Summary of the Invention - Brief Description of the Drawings - Detailed Description - Claim(s) - Abstract of the Disclosure 4. <input checked="" type="checkbox"/> Drawing(s) [total sheets <u>8</u>] 5. <input checked="" type="checkbox"/> Oath or Declaration [total sheets <u>1</u>] a. <input checked="" type="checkbox"/> Newly executed (original or copy) b. <input type="checkbox"/> Copy from prior appl. (37 C.F.R. § 1.63(d)) i. <input type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b). </div> <div style="width: 48%;"> 6. <input type="checkbox"/> Application Data Sheet. (See 37 CFR 1.76) 7. <input type="checkbox"/> CD-ROM or CD-R in duplicate, large table or Computer Program (Appendix) 8. Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary) a. <input type="checkbox"/> Computer Readable Form b. <input type="checkbox"/> Specification Sequence Listing on: i. <input type="checkbox"/> CD-ROM or CD-R (2 copies); or ii. <input type="checkbox"/> paper c. <input type="checkbox"/> Statements verifying identity of above copies </div> </div>		
<div style="display: flex; justify-content: space-between;"> <div style="width: 48%;"> 17. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment, or in an Application Data Sheet under 37 CFR 1.76: <input type="checkbox"/> Continuation <input type="checkbox"/> Divisional <input type="checkbox"/> Continuation-in-part (CIP) of prior application No.: <u> </u>/<u> </u> Prior application information: Examiner: <u> </u> Group/Art Unit: <u> </u> <small>FOR CONTINUATION OR DIVISIONAL APPS only The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 5b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.</small> </div> <div style="width: 48%;"> ACCOMPANYING APPLICATION PARTS 9. <input checked="" type="checkbox"/> Assignment Papers (coversheet/document(s)) 10. <input type="checkbox"/> 37 CFR. 3.73(b) Statement <input checked="" type="checkbox"/> Power of Attorney (when there is an assignee) 11. <input type="checkbox"/> English Translation Document 12. <input type="checkbox"/> IDS & Form 1449 <input type="checkbox"/> Copies of IDS Citations 13. <input type="checkbox"/> Preliminary Amendment 14. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) 15. <input type="checkbox"/> Certified Copy of Priority Document(s) 16. <input checked="" type="checkbox"/> Other: Certificate of Mailing by Express Mail </div> </div>		
17. CORRESPONDENCE ADDRESS		
<input type="checkbox"/> Customer Number or Bar Code Label or <input checked="" type="checkbox"/> Correspondence address below		
Name	William J. KUBIDA, Esq.	
	Hogan & Hartson, LLP	
Address	1200 17 th Street Suite 1500	
City	Denver	State CO ZIP 80202
Country	US	Telephone (719) 448-5900 Fax (719) 448-5922
Name (Print/Type)	Steven K. Barton	Registration No. 36,445
(Signature)	<i>Steven K. Barton</i>	Date 3 Nov 2000

FEE TRANSMITTAL for FY 2000

Complete if Known

Application Number	-----
Filing Date	Herewith
First Named Inventor	Richard Willems
Examiner Name	
Group / Art Unit	
Attorney Docket No.	PD99-2788

TOTAL AMOUNT OF PAYMENT (\$) **(\$)** 830.00

METHOD OF PAYMENT (check one)

1. ☐ The Commissioner is hereby authorized to charge indicated fees and credit any over payments to:

Deposit
Account
Number

50-1123

Deposit
Account
Name

Hogan & Hartson L.L.P.

- ☒ Charge Any Additional Fee Required Under 37 CFR § 1.16 and 1.17
☐ Applicant claims small entity status. See 37 CFR 1.27

2. ☒ **Payment Enclosed:**

- ☒ Check ☐ Money Order ☐ Other

FEE CALCULATION

BASIC FILING FEE

Large Entity Fee (\$)	Small Entity Fee (\$)	Fee Description	Fee Paid
710	355	Utility Filing Fee	710
320	160	Design filing fee	
490	245	Plant filing fee	
710	355	Reissue filing fee	
150	75	Provisional filing fee	

SUBTOTAL (1) **(\$)**

2. EXTRA CLAIM FEES

Total Claims	Extra Claims	Fee from below	Fee Paid
17	-20**= 0	18	0
Independent Claims 4	-3**= 1	80	80
Multiple Dependent			0

**or number previously paid, if greater, For Reissues, see below

Large Entity Fee Code	Large Entity Fee (\$)	Small Entity Fee Code	Small Entity Fee (\$)	Fee Description
103	18	203	9	Claims in excess of 20
102	80	202	40	Independent claims in excess of 3
104	270	204	135	Multiple dependent claim, if not paid
109	80	209	40	**Reissue independent claims over original patent
110	18	210	9	**Reissue claims in excess of 20 and over original patent

SUBTOTAL (2)

(\$)80.00

FEE CALCULATION (continued)

3. ADDITIONAL FEES

Large Entity Fee (\$)	Small Entity Fee (\$)	Fee Description	Fee Paid
130	65	Surcharge - late filing fee or oath	
50	25	Surcharge - late provisional filing fee or cover sheet	
130	130	Non-English specification	
2,520	2,520	For filing a request for ex parte reexamination	
920*	920*	Requesting publication of SIR prior to Examiner action	
1,840*	1,840*	Requesting publication of SIR after Examiner action	
110	55	Extension for reply within first month	
390	195	Extension for reply within second month	
890	445	Extension for reply within third month	
1,390	695	Extension for reply within fourth month	
1,890	945	Extension for reply within fifth month	
310	155	Notice of Appeal	
310	155	Filing a brief in support of an appeal	
270	135	Request for oral hearing	
1,510	1,510	Petition to institute a public use proceeding	
110	55	Petition to revive - unavoidable	
1,240	620	Petition to revive - unintentional	
1,240	620	Utility issue fee (or reissue)	
440	220	Design issue fee	
600	300	Plant issue fee	
130	130	Petitions to the Commissioner	
50	50	Petitions related to provisional applications	
240	240	Submission of Information Disclosure Stmt	
40	40	Recording each patent assignment per property (times number of properties)	40.00
710	355	Filing a submission after final rejection (37 CFR § 1.129(a))	
710	355	For each additional invention to be examined (37 CFR §1.129(b))	
710	355	Request for Continued Examination	
		Request for expedited examination of a design application	

Other fee (specify)

*Reduced by Basic Filing Fee Paid **SUBTOTAL (3)**

(\$)40.00

SUBMITTED BY

Name **Steven K. Barton**
(Print/Type)
Signature *Steven K. Barton*

Registration No. (Attorney/Agent) **36,445**

Complete (if applicable)

Telephone **(719) 448-5900**

Date **3 Nov 2000**

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of:

Richard Willems

Serial No. _____

Filed: Herewith

For: SYSTEM AND METHOD FOR COLLECTING SYSTEM
DATA USING RECORD BASED REQUESTS WITH TAG
LISTS AND PAUSING ALL BUT ONE THREAD OF A
COMPUTER SYSTEM

Group Art Unit: _____

Examiner: _____



CERTIFICATE OF MAILING BY EXPRESS MAIL

BOX PATENT APPLICATION

Assistant Commissioner of Patents and Trademarks
Washington, D.C. 20231

Sr:

The undersigned hereby certifies that the following documents:

1. Utility Patent Application Transmittal;
2. Fee Transmittal and \$ 790 filing fee;
3. Utility Patent Application;
4. Executed Declaration;
5. Executed Power of Attorney by Assignee
6. 8 sheets of drawings;
7. Recordation Form Cover Sheet PTO 1595 with Executed
Assignment and Recording Fee of \$40.00;
8. Return postcard; and
9. Certificate of Mailing By Express Mail

relating to the above application, were deposited as "Express Mail", Mailing Label No. EL700671685US
with the United States Postal Service, addressed to The Commissioner of Patents and Trademarks,
Washington, D.C., 20231, Nov 3, 2000.

Nov 3, 2000

Date

3 Nov 2000

Date

Julie M. Trout
Mailer

Steven K. Barton

Steven K. Barton, Reg. No. 36,445
Hogan & Hartson LLP
One Tabor Center
1200 17th Street, Suite 1500
Denver, Colorado 80202
Tel (719) 448-5909
Fax (303) 899-7333

**SYSTEM AND METHOD FOR COLLECTING SYSTEM DATA USING
RECORD BASED REQUESTS WITH TAG LISTS AND PAUSING ALL BUT
ONE THREAD OF A COMPUTER SYSTEM**

RELATED APPLICATIONS

DOCKET # 92090260

This application is related to a simultaneously
filed application entitled "System and Method for
5 Collecting System Data Using Automatically Identified
Symbol Libraries to Create Record Based Requests with
Tag Lists of Data to be Collected Upon Pausing All But
One Thread of an Operating Computer System",
identified as Serial No. [XX/XXX,XXX] and attorney
10 docket number PD99-2787. This application is also
related to a simultaneously filed application entitled
"Collection Driver for Collecting System Data Using
Record Based Requests with Tag Lists and Pausing All
But One Thread of a Computer System", identified as
15 Serial No. [XX/XXX,XXX] and attorney docket number
PD99-2789.

FIELD OF THE INVENTION

The invention is related to the field of
interactive debuggers for computer operating systems.
20 In particular, the invention is related to the field
of formatting requests for system information by a
debugger's user interface to control capture of list
information by the debugger's collection driver.

BACKGROUND OF THE INVENTION

Many modern computer systems, including their resident operating systems and application programs have software bugs, or may have or develop other problems that cause maloperation. When maloperation occurs, it is desirable to identify any causative bugs or other problems so that repairs may be made to prevent future maloperation. Symbolic debuggers may be used by trained service personnel to examine variables of a maloperating computer system to help those personnel identify any causative bug or other problem.

Software bugs may exist in the operating system itself or in application programs. Viruses may enter the system and cause various levels of damage, leading to maloperation. Other problems that can cause maloperation include hardware defects, malfunctions of attached devices and networked computer systems, and maliciously or accidentally incorrect user input.

Debugging is often an iterative process of information gathering, making changes or performing other experiments, and testing. A symbolic debugger is a tool highly useful for at least the information gathering phase of this process, and occasionally useful for performing experiments and testing.

Symbolic kernel debuggers are symbolic debuggers having command sets and functions particularly suited to solving software bugs and problems associated with an operating system kernel or software drivers, such as input/output device drivers that run in kernel mode.

Symbolic kernel debuggers are often used to diagnose intermittent bugs in systems. Debugging intermittent problems requires that the failure state of the intermittent bug be reproduced, then the debugger may be used to examine information relevant to the bug. Reproducing intermittent bugs can be difficult and time consuming, sometimes requiring extensive repetition of test sequences to induce the failure state.

10 **Symbolic Debuggers and Symbol Resolution**

Symbolic debuggers, including symbolic kernel debuggers, typically incorporate a collection driver, a user interface, and a symbol resolution system.

The symbol resolution system reads a module-specific symbol file, typically generated by an assembler, compiler or linker when the module was created, to obtain a list of known symbols and relative or absolute addresses corresponding to those symbols. The symbol resolution system uses this list to translate symbolic requests by service personnel into memory addresses having variables to be read or function entry points to be called or intercepted. The symbol file may also specify a variable type associated with each symbol, which may control the way the user interface displays memory data.

Symbolic kernel debuggers may have the ability to parse system information, such as process lists and input/output buffers and to display relevant information. The symbol resolution system is often used by these debuggers to locate those lists and buffers since those lists and buffers may appear at

different locations in memory for each version of the kernel or driver software.

It is essential that the version of the symbol file used by the symbol resolution system correspond to the version of the module running on the machine being debugged. If this is not the case, the debugger may read different locations than those intended, which can result in user confusion or a crashed debugger. Since operating modules are updated frequently, with service packs and in-line updates as well as with new operating system releases, locating, storing, and ensuring use of the correct symbol file can be a difficult exercise. Kernel-mode driver modules may also be released, patched, and updated, by hardware vendors; further complicating the logistical problem of ensuring use of the correct symbol file.

The collection driver typically gathers data as requested by service personnel, and, may, but need not, also have the ability to alter selected memory locations. The collection driver runs on the target machine, the machine being diagnosed. Some symbolic kernel debuggers are known to utilize a serial port of the target machine to communicate with a diagnosis machine upon which runs the user interface and symbol resolution system.

The user interface typically interfaces the collection driver and the symbol resolution system to a keyboard and display for interaction with service personnel users. The user interface may include system-specific code for reading linked-lists, including process lists, and displaying information from those lists to the user. Extraction of data from

linked lists typically requires multiple calls from the user interface to the collection driver.

Coherency of Data

When an operating system runs on a machine, it is known that many operating system variables and data structures change as the system runs. Many of these data structures are of length greater than the word length of the machine; changes made to these structures must take place over several processor operations. If these data structures are examined by a debugger after the first operation of a change, but before the last operation of the change, the data captured or viewed by the debugger may not accurately reflect the state of the system. Similarly, if a debugger begins to view or capture a data structure prior to a change, but completes capture after the change, the data captured or viewed will be incoherent in that it does not accurately reflect the state of the system.

Incoherent data may cause confusion to service personnel attempting to interpret it. Since no indication of incoherency exists, it can be difficult to determine whether a problem indicated by the data is because the data is incoherent, or because the data indicates a problem with the system. Incoherent data may cause a debugger to display erroneous information. If incoherent data is followed as part of a linked list, the debugger may crash or attempt an illegal operation. For example, if links of a doubly linked list are examined after update to the forward links, but before update to the reverse links, the reverse links are incoherent and could result in a debugger

crash if the debugger follows them. Debugger crashes may not only require that the debugger be restarted, but may require extensive work to reproduce the failure state of an intermittent bug.

5 It is desirable that captured data accurately reflect system state, or be "coherent," so that service personnel may diagnose the system without confusion, wasted effort following false leads, and without losing time to restarting crashed debuggers
10 and reproducing bugs.

Some existing symbolic kernel debuggers ignore the problem of incoherent data. A debugger believed to ignore incoherency when taking a system snapshot is Microsoft i386KD running under LiveKD by SysInternals,
15 as distributed with the book: Inside Windows 2000 3'd ed. by David Solomon and Mark Russioovich, Microsoft Press, 2000. Other debuggers such as Microsoft's i386KD without LiveKD enforce partial coherency by stopping execution of all programs, except for the
20 debugger, on the target machine until debugging is complete. Stopping execution renders the target machine temporarily unusable, disrupting any real-time control functions or network services provided by that machine. Stopping execution prevents the operating
25 system from making changes to data structures while the debugger is capturing or displaying those structures thereby preventing apparent incoherency resulting from updates to these structures as the debugger is reading them. Unless execution is stopped
30 at a time when no changes to data structures are in progress, some incoherency may, however, exist.

It can be useful to obtain a snapshot of coherent state information about a system, allow that system to continue execution for some time, and obtain a second snapshot from that system. This permits service personnel to observe how the system state changes with time, which can yield useful clues about system bugs and other problems. In particular, multiple snapshots can be useful in identifying memory and resource leaks and performance problems.

10 **Linked Process Lists**

Linked lists may be used by an operating system to store information about processes. It is known that Windows NT 4.0 and Windows 2000 store process information in a linked process list. Each node of this list may incorporate a further linked thread list as well as additional information about the process that may be useful in debugging a system. For example, in addition to pointers to a thread list, a process list node may include process name, execution priority and execution privileges.

Each node of the thread list contains list pointers to security tokens, context switches, I/O request lists, and wait blocks that can be of interest to service personnel investigating a software bug or other problem. Each node of the thread list of a process node may be linked to additional linked lists of the system.

Prior Debuggers

Microsoft Kernel Debugger i386KD (KD) is a symbolic debugger tailored for Windows NT and Windows 2000 kernel and kernel-mode driver debugging. KD is

designed for operation through a serial port of a target machine. Two machines are required, the target machine on which the system being debugged is located, having a collection driver, and an analysis machine having a symbol resolution system and a user interface. Symbol files matching the system being debugged must be present on the analysis machine. Matching symbol files are not automatically located although they are verified as correctly matching the target system. When KD is in use, all other threads on the target machine are stopped until debugging is complete, severely impacting operation of that machine. KD can, however, alter system variables and allow the system to resume operation when debugging is complete.

Statement of the Problem

Collection drivers for symbolic kernel debuggers must run with high privileges in kernel mode. Code run with those privileges poses security and bug risks, so it is desirable that it be small, with few versions.

Since the locations of kernel variables, including process lists, can vary from release to release it is desirable for a debugger to derive this information from symbol files at run time rather than embedding this information in a version-specific collection driver. It is also desirable to place other system-specific information in a command plug-in of the user interface instead of the collection driver.

It is therefore desirable to have a way of specifying system and version specific information, including list format and structure information, to be collected by a collection driver. The collection driver then interprets this specification

SUMMARY OF THE INVENTION

A network-aware, symbolic kernel debugger has been prepared. A local debugger embodiment has a collection driver, user interface, and symbol resolution system all running on the target machine. A remote debugger embodiment has the collection driver on the target machine, with the symbol resolution system and user interface on an analysis machine. A central symbol library debugger embodiment has the collection driver on the target machine, the user interface on an analysis machine, and the symbol files of the symbol resolution system on a symbol reference machine.

It has been found that most "locked-up" Windows NT and Windows 2000 systems have problems with one or more of their keyboard, mouse, and display subsystems, but can still communicate over their network port. On these systems, insufficient virtual memory can prevent users from logging in on the keyboard and display, but network communications may still operate. The remote and central symbol library embodiments take advantage of this fact by using the network port of the target machine for communications between the analysis machine and the collection driver of the target machine.

The user interface is modular. It has a common framework and command line interpreter suitable for

debugging several different operating systems on
several different platforms, with one or more command
plug-ins each suitable for debugging one specific
operating system. Each command plug-in has operating
5 system-specific knowledge and related commands. This
knowledge includes by way of example and not of
limitation knowledge of the format of each node of,
and the root symbol name of, the process list. These
commands include, for example, commands to display
10 processes, with names, and the threads linked to those
processes.

The symbol resolution system of the symbolic
kernel debugger has a library of symbol files. This
library contains symbol files suitable for several
15 operating system versions that could be found on a
target machine. The symbol resolution system
automatically identifies the operating system version
present on the target machine and locates appropriate
symbol files in its library.

20 The collection driver of the symbolic kernel
debugger is operating system specific. It is capable
of taking snapshots of the target machine system state
according to an input record list provided by the user
interface command plug-in. Coherency of each snapshot
25 is enforced through use of an operating system call to
suspend execution of all processes and threads other
than the collection driver itself for the time it
takes to capture the snapshot. After the snapshot is
collected, the suspended processes and threads are
30 restarted so that operation can continue. This
frequently permits capture of snapshots without
visible disruption of system operation.

The foregoing and other features, utilities and advantages of the invention will be apparent from the following more particular description of a preferred embodiment of the invention as illustrated in the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram of the symbolic debugger of the present invention operating with combined analysis and target machines;

Figure 2, a block diagram of the symbolic debugger of the present invention operating with separate analysis and target machines;

Figure 3, a block diagram of the symbolic debugger of the present invention operating with separate analysis and target machines and a separate symbol file library machine;

Figure 4, an illustration of the directory tree structure of the symbol file library;

Figure 5, a flowchart of the symbol library tree search performed for a symbol file when the debugger initializes for a particular target machine;

Figure 6, an illustration of the structure of process and thread lists in Windows NT & Windows 2000;

Figure 7, a block diagram of an input record list as interpreted by the collection driver of the symbolic kernel debugger;

Figure 8A, a block diagram of a memory descriptor for a scalar type that may appear in the input record list of Figure 7;

Figure 8B, a block diagram of a memory descriptor
5 for a list type that may appear in the input record list of Figure 7;

Figure 8C, a block diagram of a list element descriptor that may appear in the list memory descriptor of Figure 8B;

10 Figure 8D, a block diagram of a tag array element that may appear in the tag array of the list element descriptor of Figure 8C;

Figure 9, a block diagram illustrating the structure of captured data as placed by the collection
15 driver in the capture buffer;

Figure 10, a block diagram of a list buffer descriptor for captured data, showing list buffer element descriptors; and

Figure 11, a block diagram of a list buffer
20 element descriptor and showing how it indicates locations of node data in the capture buffer.

DETAILED DESCRIPTION OF THE EMBODIMENTS

Debugger Architecture

A first embodiment (Figure 1) of the symbolic
25 kernel debugger of the present invention is operable on a single machine, the target machine. This embodiment is useful if the software bug or other problems to be examined do not involve lockup,

freezing, or other significant problems with keyboard interface and display subsystems operating on the target machine.

In this embodiment, a collection driver 100 operates in kernel mode on the target machine. Collection driver 100 communicates 102 with a communications and collection driver interface 104 of the user interface 106 operating in user mode. The communications and collection driver interface 104 serves to interface the collection driver with a command line interpreter 108 that receives commands from service personnel, as well as formats and displays data for viewing by service personnel.

The command line interpreter 108 uses framework 110 to interpret commands and control formatting of data according to an executing command plug-in 112 selected as appropriate for the operating system kernel 114 and associated drivers 116 executing on the target machine. The executing command plug-in 112 is selected from a group of available command plug-ins 118 that may include plug-ins suitable for other operating systems; in a particular embodiment the group of available command plug-ins includes plug-ins for use with the Microsoft Windows NT 4.0 and Windows 2000 operating systems.

The command line interpreter 108 expands any symbols referenced by service personnel or by the executing command plug-in 112 through calls to a symbol manager 120. The symbol manager 120 expands these symbols by finding matching symbols in a symbol file of symbol file library 122.

5

20

30

multiple instances as target machine communications
and collection driver interface 214, analysis machine
communications and collection driver interface 210,
and as the stand-alone embodiment's communications and
5 collection driver interface 104.

This communications and collection driver
interface 214 serves to interface the collection
driver 206 with a command line interpreter 218.
Command line interpreter 218 corresponds to the
10 command line interpreter 108 (Figure 1) of the first
embodiment. Command line interpreter 218 (Figure 2)
receives commands from service personnel, as well as
formats and displays data for viewing by service
personnel on display apparatus (not shown) of the
15 analysis machine 200.

Command line interpreter 218 uses framework 220
to interpret commands and control formatting of data
according to an executing command plug-in 222
selected, as with the first embodiment, as appropriate
20 for the operating system kernel executing on the
target machine.

The command line interpreter 218 expands any
symbols referenced by service personnel or by the
executing command plug-in 222 through calls to a
25 symbol manager 224. Symbol manager 224 expands these
symbols by finding matching symbols in a symbol file
of symbol file library 226 as with the symbol library
122 of the first embodiment.

A third embodiment (Figure 3) of the symbolic
30 kernel debugger of the present invention uses a symbol
reference machine 300 as well as an analysis machine

302 and target machine 310. Each symbol reference machine 300 may serve more than one analysis machine. Symbol reference machine 300 is coupled to analysis machine (or machines) 302 through network interconnect 5 304, which may incorporate any combination of local area network, a wide area network, as well as encrypted communications over the Internet and through suitable firewalls as known in the art of virtual private networking. Encrypted communications over the 10 Internet offers particular advantage in that a symbol reference machine 300 may be maintained at a central location while the analysis machine 302 may, but need not, be a notebook computer carried by service personnel to the location of the target machine 310.

15 In this embodiment, the collection driver 306, operating in kernel mode 308, on the target machine 310 to gather data from the operating system kernel (not shown), drivers (not shown) or other programs running on the target machine 310 in manner similar to 20 the collection driver 206 (Figure 2) of the second embodiment. Collection driver 306 communicates through a local communications and collection driver interface 312 and communication service module 314 over network interconnect 316 to communications and 25 collection driver interface 318 of analysis machine 302. These components operate similar to the corresponding components of the second embodiment.

Communications and collection driver interface 318 serves to interface the collection driver 306 with 30 a command line interpreter 320, framework 322, and executing command plug-in 324 that correspond to similar components (Figure 2) of the second embodiment. Command line interpreter 320 (Figure 3)

receives commands from service personnel, as well as formats and displays data for viewing by service personnel on display apparatus (not shown) of the analysis machine 302.

5 Command line interpreter 320 expands any symbols referenced by service personnel or by the executing command plug-in 324 through calls to a symbol manager 326. Symbol manager 326 expands these symbols by seeking any matching symbols in a local symbol cache
10 328. Symbols found in the local symbol cache 328 are expanded with their definition stored therein. Symbols not found in the local symbol cache are formatted into symbol inquiry packets and transmitted over network interconnect 304 to the symbol reference
15 machine 300 for expansion.

Remote symbol manager 330 running on symbol reference machine 300 receives symbol reference packets from one or more analysis machines 302 and expands these by locating appropriate symbols in
20 appropriate symbol files of a symbol file library 332.

Symbol File Library Structure and Operation

Symbol file library 332 of the third embodiment, symbol file library 122 of the first embodiment, and symbol file library 226 of the second embodiment have
25 structure as illustrated in Figure 4. This library is stored on a memory system, such as a disk drive, of the associated machine. The library tree root 400 is a directory named SYMBOLS that appears in a known directory on a storage device, such as a disk drive or
30 RAID cluster. This directory has one or more platform subdirectories, each of which is dedicated to a

particular machine architecture and operating system family. For example and not by way of limitation, there may be a platform subdirectory named INTEL_NT 402 for Windows NT and similar operating systems
 5 operating on Intel Pentium type machines, a platform subdirectory named Intel_95 404 for Windows 96, Windows 98, and Windows Millennium Edition operating on Intel Pentium machines, and additional platform directories 406 for other machine-system family
 10 combinations. Additional platform directories 406 may include directories for Linux running on Intel machines, for Windows NT running on future Intel 64-bit machines, and other combinations.

Within a platform directory, such as the Intel_NT
 15 directory 402, there are build number subdirectories for each released build. For example, Windows NT 4.0 has build number "1381", so files related to this version are placed in a subdirectory named "1381" 408. Similarly, Windows "2000" has build number "2195", so
 20 files related to this version are placed in a subdirectory named "2195" 410.

Within each build number subdirectory is one or more additional service pack subdirectories, one directory corresponding to each released service pack
 25 of the system, with a zero directory for the initial release. For example, service pack subdirectory "0" 412 and service pack subdirectory "1" 414 may exist in build subdirectory "2195" 410. Within each of these service pack subdirectories is another subdirectory, a
 30 service pack symbols directory, named SYMBOLS, such as the service pack symbols directory for service pack 0 of Windows "2000" 416.

Within each service pack symbols directory, such as service pack symbols directory 416, are three symbol file subdirectories named "EXE" 418, "DLL" 420, and "SYS" 422. The "EXE" subdirectory contains symbol files associated with kernel mode and driver executable files ending in the ".EXE" suffix, the "DLL" subdirectory contains symbol files associated with kernel mode and driver executable files ending in the ".DLL" suffix, and the "SYS" subdirectory contains symbol files associated with kernel mode and driver executable files ending in the ".SYS" suffix. These symbol files may include ".DBG" and ".PDB" files generated when corresponding executable files are compiled, such as symbol files 424 and 426.

Within the service pack subdirectory, such as service pack subdirectory "0" 412, there is also a subdirectory named HOTFIXES 430 that may contain one or more hotfix subdirectories corresponding to hotfixes or in-line fixes applicable to a particular service pack of a system release. Each hotfix subdirectory, such as hotfix subdirectory Q123456 432, has a name corresponding to the hotfix release. Within the hotfix subdirectory are three symbol file subdirectories named "EXE" 434, "DLL" 436, and "SYS" 438. The "EXE" subdirectory contains symbol files associated with kernel mode and driver executable files ending in the ".EXE" suffix, the "DLL" subdirectory contains symbol files associated with kernel mode and driver executable files ending in the ".DLL" suffix, and the "SYS" subdirectory contains symbol files associated with kernel mode and driver executable files ending in the ".SYS" suffix. These symbol files may include ".DBG" and ".PDB" files

generated when corresponding executable files are
compiled, and made available by the system vendor,
such as symbol files 440 and 442. ".DBG" files are
often available from Microsoft for use with both
5 Windows NT and Windows 2000, while ".PDB" files are
available for Windows 2000 but not for Windows NT.

When the an embodiment as illustrated in Figure
1, 2, or 3, initializes for a particular target
machine, any required network connection to the target
10 machine is set up, the collection driver and
communication service if required for that platform is
started, and user interface 106, 216, or 340 of the
debugger connects 500 (Figure 5) through to the
collection driver. If connection was unsuccessful, an
15 error is reported 502 and operation ceases. If
connection was successful, the build number and
service pack number, or other version identifying
information, is obtained 504 by the collection driver
from the operating system of the target machine and
20 returned to the user interface. A list of loaded
executable modules on the target machine is also
obtained 506, this list includes module names and
checksums of the corresponding executable files. The
build number and service pack number, or other version
25 identifying information of the operating system, is
used to ensure a compatible command plug-in 112, 222,
or 324 is loaded, and used to locate symbol files in
the library tree structure.

Once the version identifying information is read
30 from the system running on the target machine, the
symbol file for the operating system kernel is
searched for in the symbol library structure. For
example, if the platform is Intel_NT, build number is

"2195", and service pack number is "0", symbol manager 120 or 224, or remote symbol manager 330, finds the SYMBOLS subdirectory 416 (Figure 4) of service pack directory 412 of build directory 410 of platform
 5 directory 402 of the library tree root 400.

The operating system kernel file name is known to the symbol manager 120, 224, or 326 (for Windows NT and Windows 2000 this name depends on the number of
 10 processors in the machine). This filename is set 507 as the symbol file to locate and a symbol file of this name having a symbol file suffix is sought 508 in the appropriate "EXE" 418, "SYS" 422, or "DLL" 420 subdirectory.

Each ".DBG" or ".PDB" symbol file has an
 15 indicator of the checksum of the corresponding executable file. These checksums for hotfix executables are entered into a small database of symbol file names and locations that can be accessed by module name, build number, service pack number, and
 20 checksum. If no symbol file is found 510, then this database is checked 512 for a matching file. If a kernel symbol file is found at this point, it is a hotfix symbol file located somewhere under the HOTFIXES directory 430. If 514 no symbol file is
 25 found, and the file being sought is the kernel symbol file then the user interface disconnects 516 from the target machine and reports an error 518.

If a symbol file was found, the checksum indicated in the symbol file is compared 520 with the
 30 checksum for that module obtained with the list of loaded modules on the target machine. If 522 these checksums match, success is declared 524, the symbol

file is loaded, and the user interface awaits a user command.

If the checksums do not match, the symbol file hotfix index database is checked 522 for a symbol file
5 having the same module name, build, service pack, and checksum, and the associated symbol file is read. This symbol file is located in an appropriate "EXE" 434, "SYS" 438, or "DLL" 436 subdirectory of a hotfix directory. If 514 one is found, success is declared
10 522, the symbol file is loaded, and the user interface awaits a user command.

As debugging progresses, it may be found necessary to locate additional symbol files corresponding to other modules. This is done by
15 reentering 530 the above sequence to search for a symbol file having an appropriate filename, including 508, 510, 512, 514, 518, 520, 522, and 524 for those modules.

Most hotfix releases include updated versions of
20 some, but not all, operating system executables and corresponding symbol files. The symbol file library stores complete copies of all operating system symbol files for each build and service pack. It also stores symbol files updated in each hotfix. Those symbol
25 files corresponding to modules not updated in a particular hotfix are located in the complete build and service pack directories. The symbol library search method of the present invention will locate any matching file, whether it be in a hotfix directory or
30 in a build and service pack directory.

If more than one analysis machines 302 (each coupled to a target machine) are in use and connected to a single remote symbol manager 330, remote symbol manager 330 maintains a separate list of opened symbol files for each analysis and target machine combination. This permits the remote symbol manager 330 to respond to expansion requests with symbols appropriate to the operating system version, service pack, or hotfix, installed on each target machine.

10 **Collection Driver Input Record List**

The input record list specifies information to be captured by the collection driver.

When the command line interpreter 108, 218, or 320, requires data from the target machine, it creates an input record list and transmits this list to the collection driver. The input record list specifies information for the collection driver to gather from the operating system on the target machine. The collection driver gathers the desired information according to this list, places the information in the capture buffer, and copies the capture buffer to the command line interpreter.

System Data Structures to be Captured

It is known that much of the system data of interest to service personnel debugging a Windows 2000 or Windows NT is in linked lists. Additional system data of interest may be in scalar variables.

A particular symbol 600 (Figure 6) can be resolved by the symbol resolution system to determine an address of the process list head 602 in target

machine memory. This list head 602 contains pointers to links 604 in each process node, such as process node 606; these lists are doubly linked such that links include a forward and reverse link. Links 604 may point to links 608 in other process nodes 610 and 612 as known in the art. Process nodes contain process identity information 614, priority information 616, and other information that can be of interest in debugging.

Process nodes, such as process nodes 606 and 610, may contain a thread list head 618 containing pointers to links 620 in a list of thread (ETHREAD in Microsoft parlance) nodes such as node 622. For example and not by way of limitation, note that thread nodes 622 and 624 are linked to process node 606, and thread node 626 are linked to process node 610. Each thread node corresponds to a thread that exists in the system and may contain information of interest in debugging such as thread identity 630, wait block list head 632, and an I/O request list head 634.

Input Record List Structure

The input record list has a header 700 (Figure 7), a count of records in the list 702, one or more record offsets 704 and 706 for locating records in the list, and one or more records 708 and 710. Each record, such as record 708, of the input record list begins with a record signature 712, which is a constant that can be used to confirm the starting point of a record. Each record also has a descriptor type 714 that indicates the type of a following memory descriptor 716.

The memory descriptor 716 may be of several types as indicated by the descriptor type 714, including a scalar memory descriptor and a list memory descriptor. A scalar memory descriptor contains an address 800 (Figure 8A) and a length 802 of data to be captured. This scalar memory descriptor type is useful for capturing scalar values located at symbolically locatable locations in memory of the target machine.

The list memory descriptor has a count of list element descriptors 820 (Figure 8B), followed by one or more list element descriptor offsets 822 and 824. Each list element descriptor offset may be used to locate an associated list element descriptor, offset 822 indicates a location of list element descriptor 826 and offset 824 indicates a location of list element descriptor 828.

Each list element descriptor has a list type 840 (Figure 8C), which indicates a type of linkage used by the list having data to be captured, which may have the structure of Figure 6. For example, the list type 840 may indicate whether the list is singly linked, doubly linked, or singly linked with an element count.

The list element descriptor also has a link offset 841 that indicates the position of the links associated with the particular linked list in a node. Further, this offset permits list constructs, such as used in Microsoft operating systems, that have links at locations other than at the beginning of each node.

Each list element descriptor also has a list head type 842, which indicates whether its list head locator 844 represents a location in system memory or

an offset into a node of a parent list. It also has a node count limit 846 indicating a maximum number of list nodes from which data will be captured by the collection driver, a tag count 848 indicating how many portions of each node will be captured, and one or more tags in a tag array 850. There is also a list identifier 852 that is copied into the capture buffer to identify data captured according to a particular list element descriptor.

The collection driver can follow the list head locator 844 to find the list head in target system memory. The driver can then follow the list to each node in succession by interpreting links of the list according to the list head type 842. Once a node of the list is located, the driver can interpret the tags of the tag array 850 to identify the data to be captured from that node. Similarly, the driver can locate the heads of additional lists in the node as specified by subsequent list element descriptors, follow those heads to nodes of those lists, and capture data from those nodes.

Each tag has an offset 870 (Figure 8D) and a length 872 indicating a portion of a list node to be captured.

Input Record List Operation

Assume by way of example and not of limitation that the user interface wants a snapshot comprising a thread count from a scalar variable, and thread and process identification information from a Windows NT or Windows 2000 process list.

The user interface, including the command plug-in, will cause construction of an input record list containing at least three records. With reference to Figures 7, 8A, 8B, 8C, and 8D, a first record 708 of scalar memory descriptor type is placed in the input record list specifying an address 800 and length 802 of the thread count scalar variable. A second record 710 is constructed having a list memory descriptor type 730 and suitable memory descriptor 732.

The memory descriptor 732 for this operation specifies that it has two list element descriptors 826 and 828. The first of these list element descriptors 826 is used to specify the process list, the second of these list element descriptors 828 is used to specify the thread list appurtenant to the process list.

With reference to Figures 6, 8B, 8C, and 8D, list element descriptor 826 list head type 842 is marked to indicate that the list head locator 844 specifies the location in system memory of the process list head 602. An entry is made in the tag array 850 to specify an offset 870 and length 872 of a tag corresponding to the process identifier location 614 in each process node. A node count limit 846 is set to prevent capture buffer overflow.

The second list element descriptor 828 is marked to indicate that its list head locator 844 specifies a location in a node of a parent list instead of a node in memory. The list head locator 844 specifies the offset of the thread list head 618 in a process node, such as process node 616. The node count limit 846 is set to a reasonable value to prevent freezing of the system should a circular list be found, and a tag

array 850 entry is made having the offset 870 and
length 872 of the thread identification 630 of a
thread list node, such as node 622. Offset 870 of the
thread list head is derived from information in header
5 files compiled into the command plug-in of the user
interface. The locations of the process list header
844 is derived from information in an associated
symbol file previously extracted from the symbol file
library.

10 Once assembled, the input record list is
transmitted to the collection driver. The collection
driver 100, 206, or 306, then uses a system call to
briefly suspend execution of all threads running on
the target machine except itself, interprets the input
15 record list while copying a snapshot of system data to
the capture buffer, and uses another system call to
restart execution of the suspended threads.

On a single processor machine, it may be
sufficient to disable processor interrupts before
20 capturing the system information, re-enabling
interrupts afterwards. On multiple processor machines
having shared memory it is necessary to temporarily
stop execution of threads on additional processors.
The collection driver as built for Windows NT and
25 Windows 2000 suspends execution of all other threads
with the KeContextToKframes and KeFreezeExecution
system calls, and resumes execution with the
KeThawExecution call.

This capture process takes about twenty
30 milliseconds on a 400 Mhz Pentium-class machine for a
typical snapshot of the many elements of system data
required by a practical symbolic kernel debugger.

Since execution is resumed following the brief capture process, the target machine may continue operation. The capture process is brief enough that many, but not all, real-time control programs and drivers may be
 5 debugged without disrupting the system being controlled.

Capture Buffer Format

System data gathered by the collection driver is placed in the capture buffer with structure as
 10 illustrated in Figure 9. The buffer contains a count 900 of the number of records in the buffer, and offsets 902 and 904 indicating locations of records in the buffer. With reference to Figure 7, 8A, and 9, each record, such as records 906 and 908, contains a
 15 record signature 910, a result flag 912, a capture type 914, a buffer descriptor 916, and data 918.

When data is captured according to a record of scalar data descriptor type, the capture type 914 of the corresponding record is marked with a scalar
 20 buffer descriptor type and the buffer descriptor 916 is set to the address 800 and length 802 of the data 918 in the buffer.

Similarly, when data is captured according to a record of list descriptor type, the capture type 914
 25 of the corresponding record is marked with a list buffer descriptor type.

If the capture type 914 is list buffer descriptor type, the buffer descriptor 916 has structure as illustrated in Figure 10. It comprises a count 1000
 30 of list buffer element descriptors and one or more list buffer element descriptor offsets 1002. Each

list buffer element descriptor indicates the location of a list buffer element descriptor, such as element descriptors 1004 and 1006, in the buffer.

Each element descriptor, such as element
5 descriptor 1004, in turn has structure according to
Figure 11. Each list buffer element descriptor has
the number of nodes captured 1100 at this level of the
list, the size of each node 1102 as captured (after
the requested tag operations limit the amount of
10 information captured from the node), a result flag
1104 indicating whether capture was successful, and a
data offset 1106 indicating the location of the first
captured node data in the record.

Summary

15 The collection driver therefore is capable of
capturing a coherent snapshot of critical system data
in the capture buffer based on the input record list.
Once captured, the capture buffer is dispatched to the
user interface for interpretation, formatting, and
20 display.

The embodiments have been described with particular reference to a symbolic kernel debugger for debugging software bugs and other problems with the Microsoft Windows NT and Windows 2000 operating systems. The principles of the kernel debugger can be utilized in symbolic kernel debuggers for debugging on other operating systems including by way of example the Microsoft Windows 98 and Windows Millennium Edition, as well as operating systems from other vendors including Unix and Unix-derived operating systems such as Solaris, Tru-64 Unix, and Linux.

The embodiments have been described with reference to a particular symbol file library tree structure. It is anticipated that the invention will be operable with other symbol file library structures searchable by operating system version information, such a build number and service pack number.

It is also anticipated that the invention will be operable with alternative ways of verifying that a located symbol file matches a loaded module of the operating system. Verification in the described embodiments is done with the checksum matching process discussed above, the checksum serving as file tagging information. Verification can also be accomplished through comparison of alternative tagging information such as a timestamp, a hash function of corresponding source code, or both, should these tags be incorporated into the executable and symbol files.

A presently preferred embodiment of the present invention and many of its improvements have been described with a degree of particularity. It should be understood that this description has been made by way of example, and that the invention is defined by the scope of the following claims.

A computer program product is any machine-readable media, such as an EPROM, ROM, RAM, DRAM, disk memory, or tape, having recorded on it computer readable code that, when read by and executed on a computer, instructs that computer to perform a particular function or sequence of functions. A computer having the code loaded on it includes a computer program product because it incorporates DRAM and/or disk memory having the code recorded in it. It

A presently preferred embodiment of the present invention and many of its improvements have been described with a degree of particularity. It should be understood that this description has been made by way of example, and that the invention is defined by the scope of the following claims.

CLAIMS

What is claimed is:

- 1 1. A method for gathering data from memory of a
2 computer system, comprising the steps of:
3 following a plurality of memory element
4 descriptors of a machine readable record list to
5 locate data in the memory of the computer system,
6 where each memory descriptor is descriptive of data to
7 be retrieved from memory of the computer system;
8 gathering data specified by the plurality of
9 memory element descriptors; and
10 formatting the data into a buffer.
- 1 2. The method for gathering data from memory of a
2 computer system of Claim 1, wherein at least one
3 memory descriptor is descriptive of a list memory
4 type, including location information of a head of a
5 list and tag information for at least one data element
6 to be gathered from a node of the list.
- 1 3. The method for gathering data from memory of a
2 computer system of Claim 1, wherein at least one
3 memory descriptor is descriptive of a scalar memory
4 type.
- 1 4. The method for gathering data from memory of a
2 computer system of Claim 3, wherein at least one
3 memory descriptor is descriptive of a list memory
4 type, including location information of a head of a
5 list and tag information for at least one data element
6 to be gathered from a node of the list.

1 5. The method for gathering data from memory of a
2 computer system of Claim 3, wherein at least one
3 memory descriptor is a list memory descriptor,
4 including location information of a head of a first
5 list, location information of a head of a second list
6 in nodes of the first list, and tag information for at
7 least one data element to be gathered from nodes of
8 the second list.

1 6. A method for parsing a linked list to extract
2 data therefrom, the linked list stored in memory of a
3 computer system, comprising the steps of:

4 constructing a record list, the record list
5 comprising at least a first list element descriptor
6 descriptive of data to be retrieved from a first
7 linked list;

8 following a list head locator of the list element
9 descriptor to a head of the first linked list;

10 following links of the head of the first linked
11 list to a first node of the linked list;

12 interpreting at least one tag of the first list
13 element descriptor to locate data of the node; and

14 extracting data from the node.

1 7. The method of parsing a linked list of Claim 6,
2 wherein:

3 the record list further comprises a second list
4 element descriptor descriptive of data to be retrieved
5 from a second linked list, and wherein a node of the
6 first linked list contains a head of the second linked
7 list; and

8 the method further comprises the steps of:

9 following a list head locator of the second list
10 element descriptor to a second list head of the node
11 of the first linked list;

12 following links of the second list head to a node
13 of the second list;

14 interpreting at least one tag of the second list
15 element descriptor to locate data of the node of the
16 second list; and

17 extracting data from the node of the second list

1 8. The method of parsing a linked list of Claim 7,
2 further comprising the step of formatting the
3 extracted data into a capture buffer.

1 9. The method of parsing a linked list of Claim 7
2 further comprising the steps of
3 stopping execution of all threads executing on
4 the computer system except for a thread parsing the
5 list; and

6 resuming execution of all threads stopped during
7 the step of stopping execution;

8 wherein the step of stopping execution is
9 performed prior to the step of following links of the
10 head of the first linked list, and the step of
11 resuming execution is performed after the step of
12 extracting data from the node of the second list.

18 symbol resolution system, and transmitting the input
19 record list to the collection driver;

20 wherein the collection driver further comprises
21 code for interpreting the input record list and
22 collecting operating system data into a capture buffer
23 specified by the input record list, and transmitting
24 the capture buffer to the user interface.

1 14. The symbolic debugger of Claim 13, wherein the
2 collection driver is capable of interpreting a record
3 of the input record list that specifies information to
4 be gathered from multiple nodes of a linked list.

1 15. The symbolic debugger of Claim 14, wherein the
2 collection driver is capable of interpreting a record
3 of the input record list that specifies information to
4 be gathered from multiple nodes of a linked list
5 having a list head located in a node of a parent list,
6 a list head of the parent list being specified by a
7 record of the input record list.

1 16. The symbolic debugger of Claim 14, wherein the
2 collection driver is capable of interpreting a record
3 of the input record list that specifies scalar
4 information to be gathered from designated locations
5 of the memory system.

1 17. The symbolic debugger of Claim 14, wherein the
2 collection driver further comprises a communications
3 interface capable of receiving the record list over a
4 network connection and comprises computer readable
5 code for reading the version information from the
6 operating system executing on the target machine.

A method for gathering data from memory of a computer system is operable to gather scalar information, linked list information, or both. When gathering data from a linked list the method involves parsing a linked list to extract data therefrom by constructing a record list having a list memory element descriptor descriptive of the list head and data to be retrieved the list, following a list head locator of the list element descriptor to a head of the linked list, following links of the head and nodes of the list to a node of the linked list, interpreting a tag of the list element descriptor to locate data of the node, and extracting data from the node.

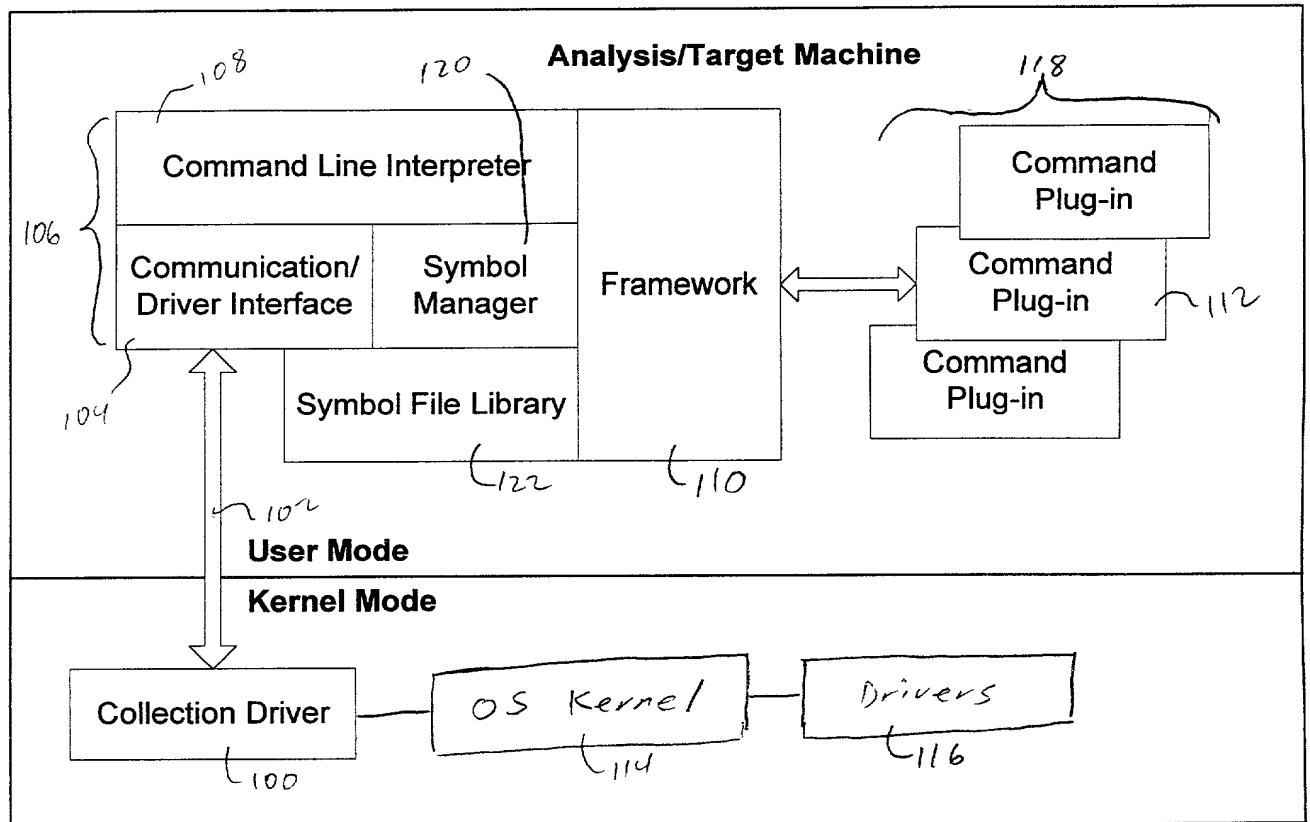


Figure 1, Local




Figure 2, Remote

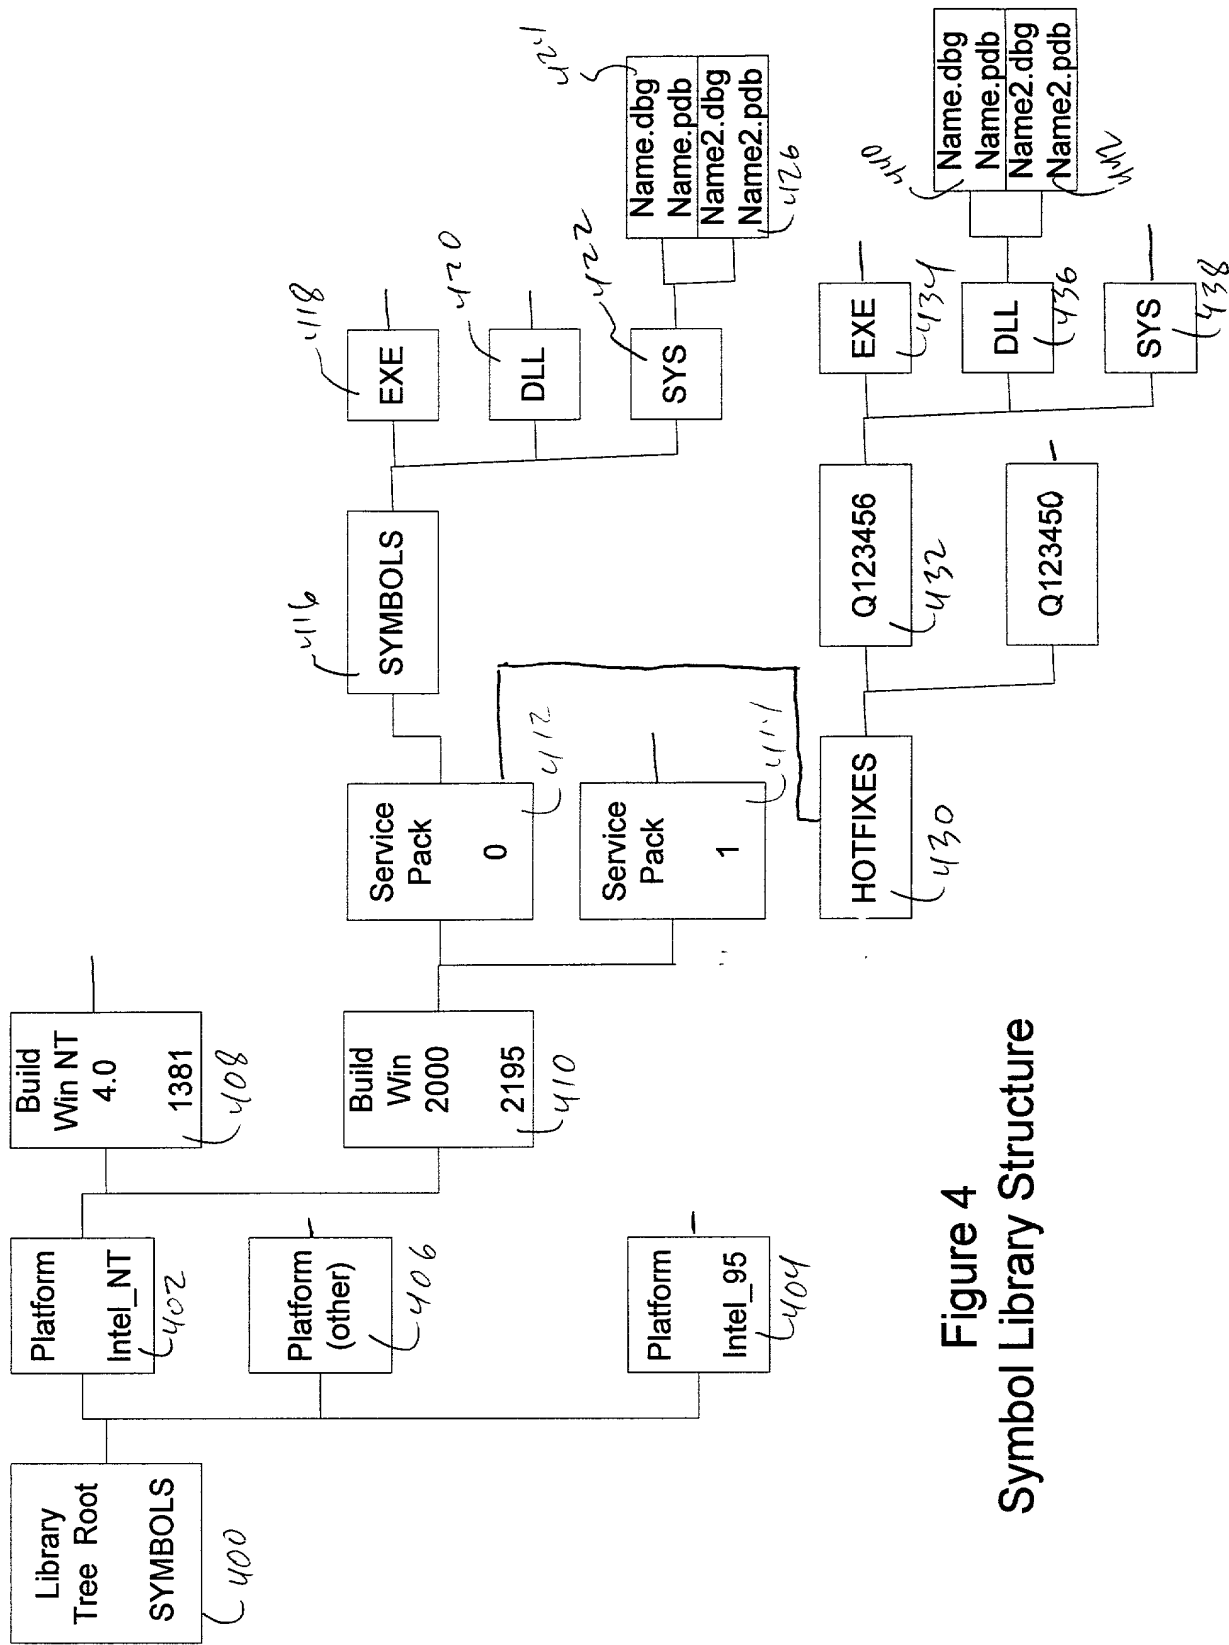


Figure 4
Symbol Library Structure

09706076 340300

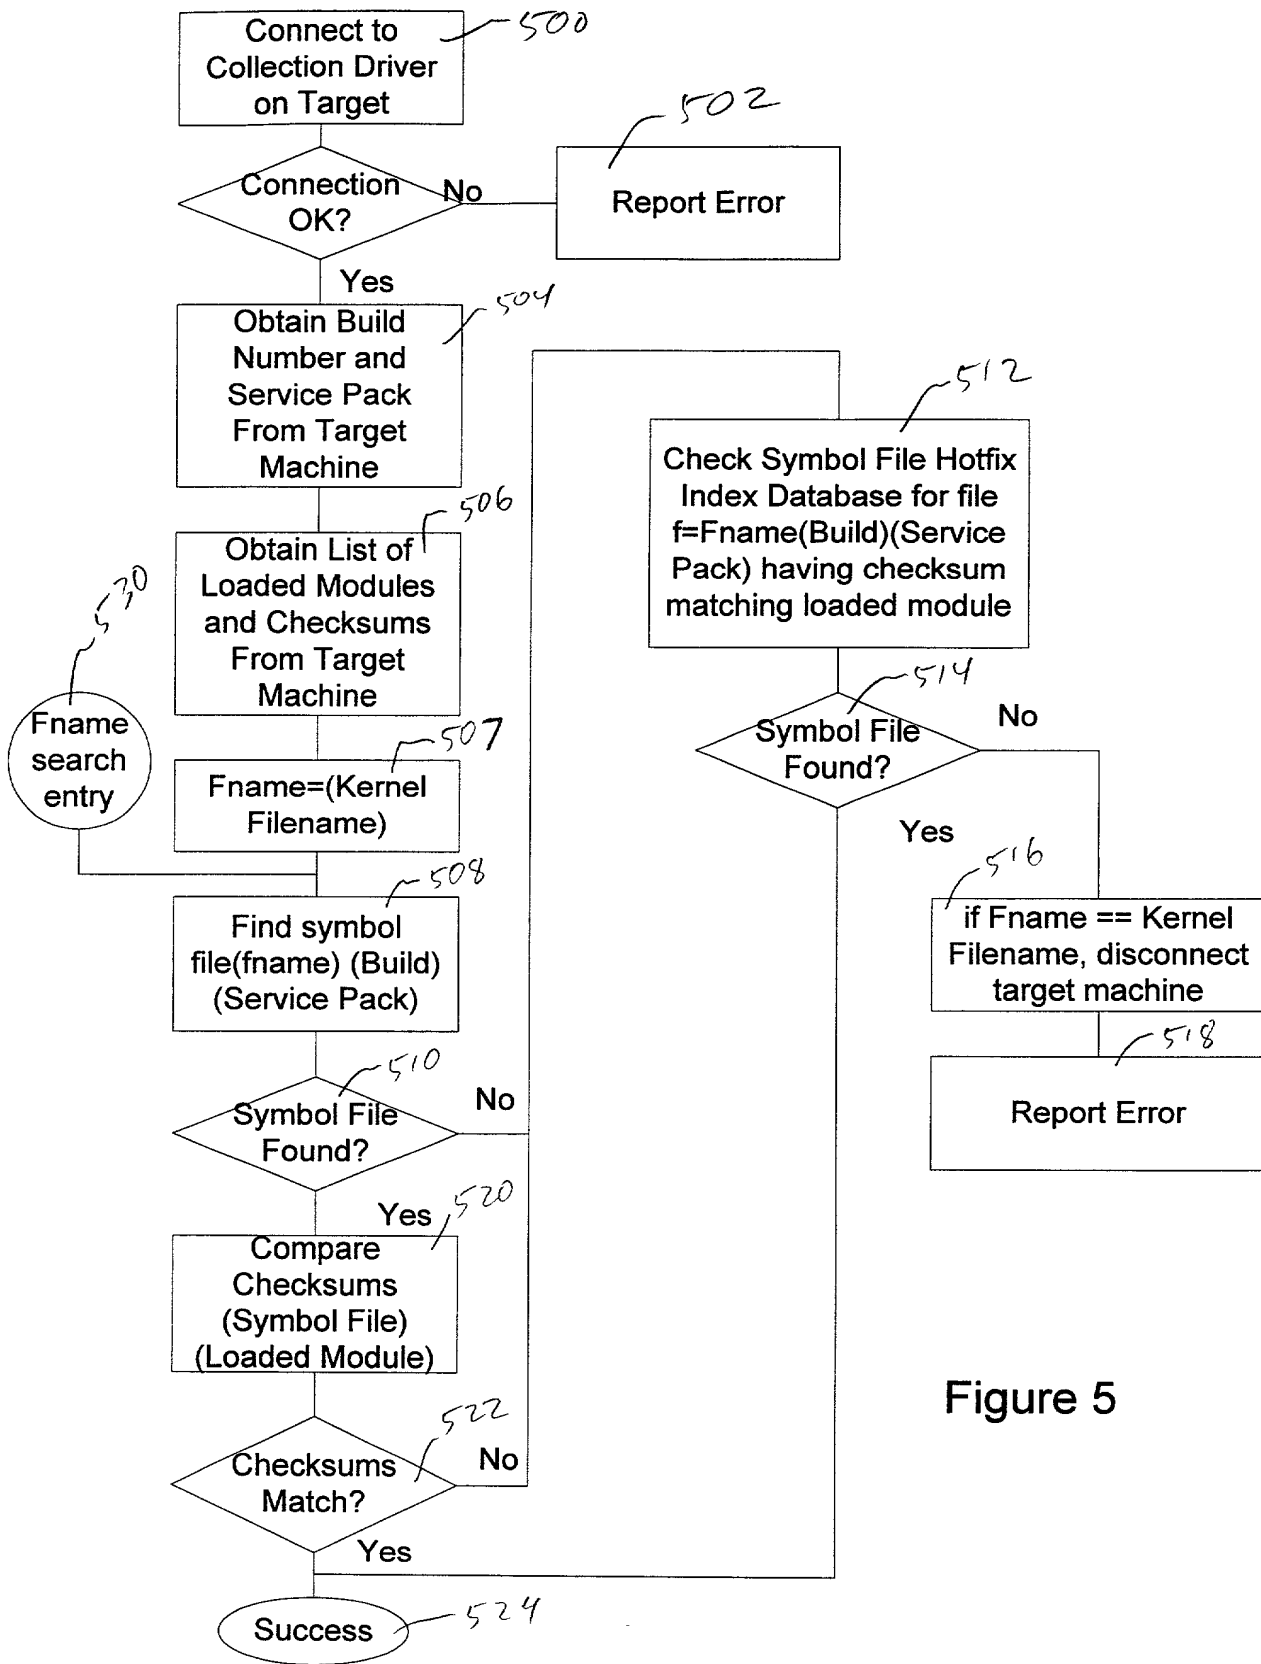


Figure 5

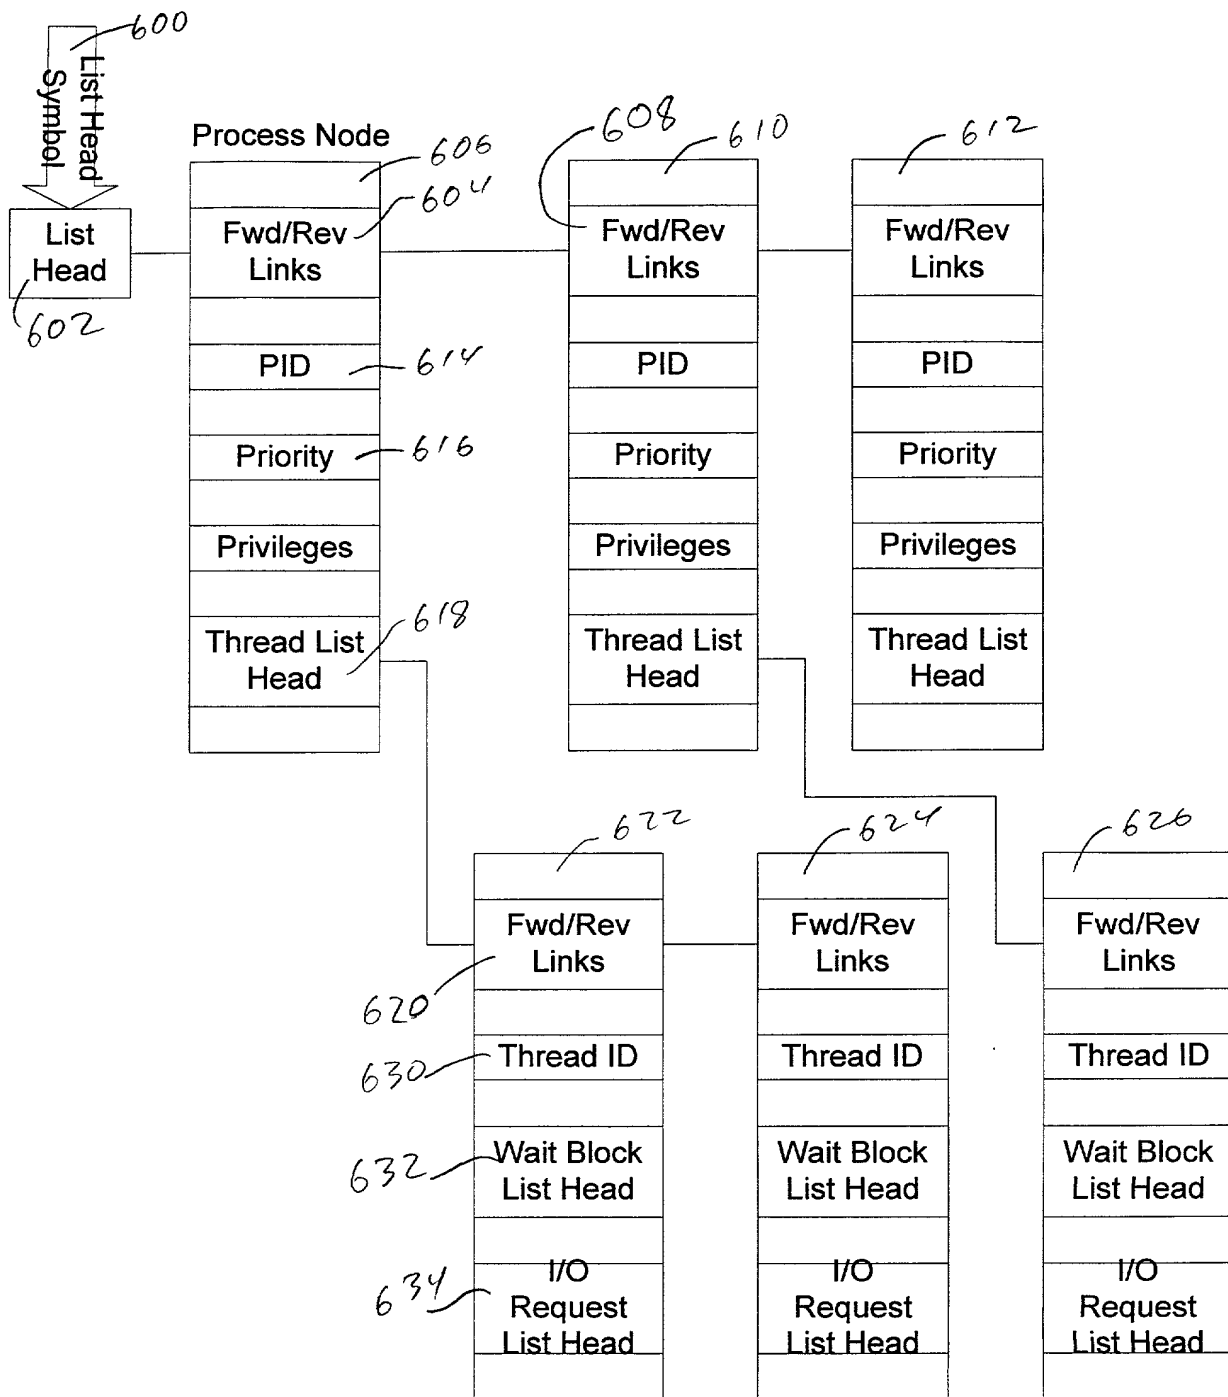


Figure 6
 KNOWN ART
 Windows NT/2000 Process List Structure




Figure 7
Input Record List Structure




Figure 8D
Tag




Figure 8A
Scalar Memory Descriptor

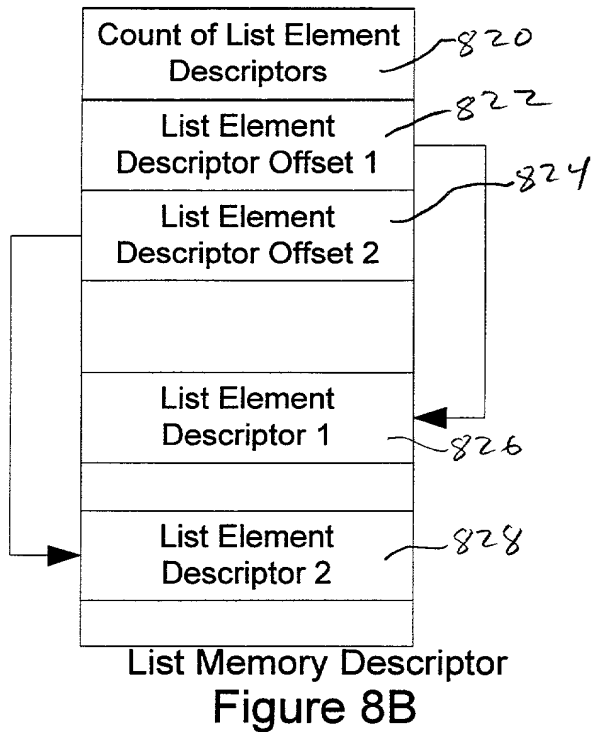


Figure 8B
List Memory Descriptor

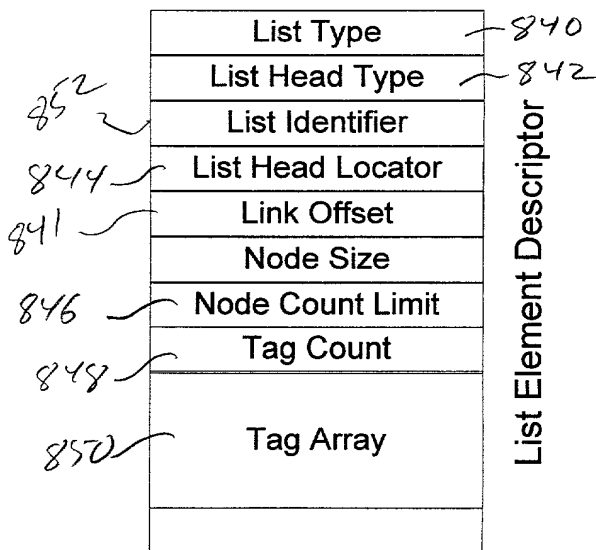


Figure 8C

09706076-11000




Figure 9
Capture Buffer Structure

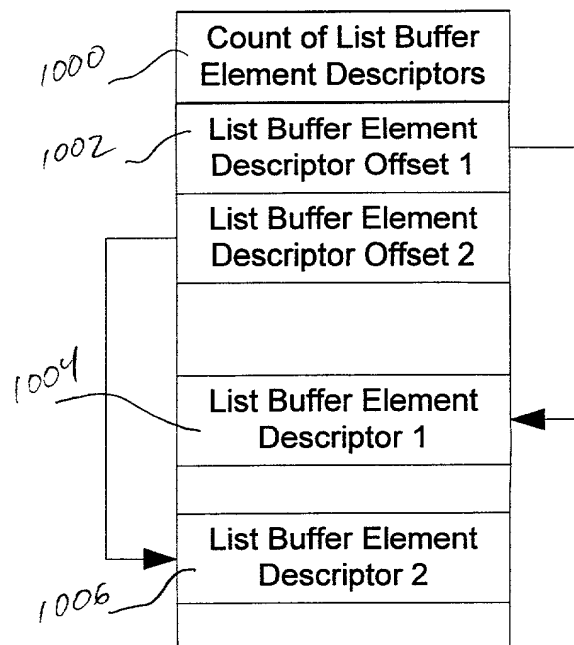


Figure 10
List Buffer Descriptor

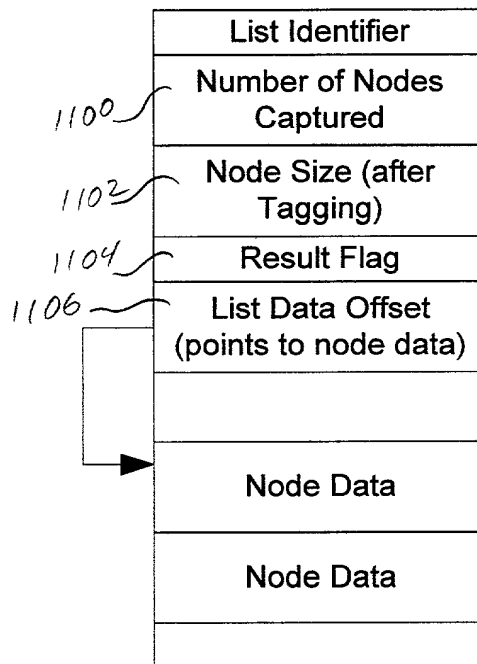


Figure 11
List Buffer Element Descriptor

DECLARATION

As a below named inventor, I hereby declare that: my residence, post office address, and citizenship are as stated below next to my name. I believe I am the original, first, and sole inventor (if only one name is listed below) or a joint inventor (if plural inventors are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

SYSTEM AND METHOD FOR COLLECTING SYSTEM DATA USING RECORD BASED REQUESTS WITH TAG LISTS AND PAUSING ALL BUT ONE THREAD OF A COMPUTER SYSTEM

as described in the specification ☒ attached or ☐ of patent Application Serial No. -----
filed ----- and amended on -----.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above; that I do not know and do not believe the same was ever known or used in the United States of America before my or our invention thereof, or patented or described in any printed publication in any country before my or our invention thereof or more than one year prior to this application; that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representative or assigns more than twelve months prior to this application; and that I acknowledge the duty to disclose information of which I am aware which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations ' 1.56(a). Such information is material when it is not cumulative to information already of record or being made of record in the application, and


- (1) it establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) it refutes, or is inconsistent with, a position the applicant has taken or may take in:
 - (i) opposing an argument of unpatentability relied on by the Office, or
 - (ii) asserting an argument of unpatentability.

I hereby claim foreign priority benefits under Title 35, United States Code ' 119 of any foreign application(s) for patent or inventor's certificates listed below and have also identified below any foreign application(s) having a filing date before that of the applications(s) on which priority is claimed:

COUNTRY	APPLICATION NUMBER	Date Filed	Priority Claimed under 35 USC 119
			<input type="checkbox"/> YES <input type="checkbox"/> NO

I hereby claim the benefit under Title 35 United States Code ' 120 of any United States application(s) listed below and, insofar as any subject matter of any claim of this application is not disclosed in the prior United States Application, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations ' 1.56(a) which occurred between the filing date of the prior application and the national PCT international filing date of this application.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

FULL NAME OF SOLE OR FIRST INVENTOR Richard A. Willems	INVENTOR'S SIGNATURE 	DATE 11-2-2000
RESIDENCE 4840 Teton Place Colorado Springs, CO 80918		CITIZENSHIP USA
POST OFFICE ADDRESS Same		
FULL NAME OF SECOND JOINT INVENTOR	INVENTOR'S SIGNATURE	DATE
RESIDENCE		CITIZENSHIP
POST OFFICE ADDRESS Same		
FULL NAME OF THIRD JOINT INVENTOR	INVENTOR'S SIGNATURE	DATE
RESIDENCE		CITIZENSHIP
POST OFFICE ADDRESS Same		

NOV-02-00 16:48 From:HOGAN&HARTSON

7194485922

T-470 P.05/11 Job-223

Express Mail No.: EL700671685US

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant/Patentee:

Richard Willems

Serial No.: -----

Date Filed: Herewith

Attorney File No.: 68854.0157

Digital Docket No.: PD99-2788

For: SYSTEM AND METHOD FOR COLLECTING
SYSTEM DATA USING RECORD BASED REQUESTS
WITH TAG LISTS AND PAUSING ALL BUT ONE
THREAD OF A COMPUTER SYSTEM

POWER OF ATTORNEY BY ASSIGNEE

Under the provisions of 37 C.F.R. § 3.71, the undersigned assignee of record of the entire interest in the above-identified patent/patent application by virtue of an assignment recorded (check as applicable):

☒ Concurrently Herewith
☐ Date Recorded _____
☐ Reel _____ Frame _____

elects to conduct the prosecution of the application/maintenance of the patent to the exclusion of the inventor(s). The undersigned hereby declares that she has reviewed the above-referenced assignment and hereby declares that, to the best of her knowledge, title is in the Assignee, and further declares that all statements made herein of her own knowledge are true and that all statements made on information and belief are believed to be true. The assignee hereby revokes any previous powers of attorney and appoints the following to prosecute this application/maintain this patent and transact all business in the Patent and Trademark Office connected therewith:

(Prosecuting Attorney List)

Irene Kosturakis, Reg. No. 33,724	Stuart T. Langley, Reg. No. 33,940
Rich Lange, Reg. No. 27,286	Carol W. Burton, Reg. No. 35,485
Louis Brucculeri, Reg. No. 38,834	Steven Kent Barton, Reg. No. 36,445
Sarah T. Harris, Reg. No. 35,891	E. Michael Byorick, Reg. No. 34,131
Joseph Arrambide, Reg. No. 39,589	Matthew G. Dyor, Reg. No. 42,278
Kelth Lutsch, Reg. No. 31,851	Steven C. Petersen, Reg. No. 36,238
Theodore S. Park, Reg. No. 26,971	Sarah S. O'Rourke, Reg. No. 41,226
William J. Kubida, Reg. No. 29,664	Kent A. Lembke, Reg. No. 44,866

Please direct all communications relative to this application to the following addressee:

WILLIAM J. KUBIDA
Hogan & Hartson LLP
One Tabor Center
1200 17th Street, Suite 1500
Denver, Colorado 80202
(719) 448-5909

ASSIGNEE

COMPAQ COMPUTER CORPORATION

Date: 02 Nov 2000BY: Diane H. Strong

NAME: Diane H. Strong

TITLE: Administrator, Patents

Authorized To Sign This Document On Behalf Of Compaq Computer Corp.